DISOC-CII (Mr. Haskins, Ms. Steffens, Mr. Given, Mr. Engel)

SUBJECT: Trip Report for Attendance at BMC Users Group and CMG 2002 International Conference.

TO: DSIO-M, Bruce Jarvis

1. AUTHORITY AND ITINERARY:

Travel Order Nos. DSIO-C-UXJ13SO 4040, 4041, 4042, and 4045 from Columbus, Ohio to Reno, Nevada and return.

2. PURPOSE OF VISIT:

To attend the BMC annual Users Group Meeting and the 2002 Computer Measurement Group International Conference.

3. PERSONS CONTACTED:

N/A

4. DISCUSSIONS AND OBSERVATIONS:

a. Sunday workshops on 8 Dec. 2002 Reno, Nevada.

1. Mark Friedman gave a Sunday morning half-day workshop on "Windows Net: An Overview of Performance." Windows .Net is the next version of Windows after Windows XP. .Net uses the XP kernel and provides new Web application runtime services. .Net was more designed for application programmers than performance tuners (as it has few system-level tuning knobs). It is difficult to support a diverse workload mix without allowing a lot of extra capacity as these applications are "brittle" when you run several on the same server; so the recommended utilization level is running the server 30-40% busy. While under Windows 2000 backups often cannot be restored (as some files were open during the backup), under .Net restoring backups is expected to work a lot better. The extra cost Windows .Net Resource Kit ($59.99 list price) is a required tool for performance analysis. The performance metrics express the queue length in threads in .Net instead of using processes like the UNIX world does. The new virtual memory manager in .Net will allow applications to use 3 GB of memory instead of 2 GB if needed.

2. Randy Kerns of Evaluator Group Inc. gave a Sunday afternoon half-day workshop on "Storage Management in a Storage Network." Storage networks can allow a

four-fold increase in the amount of storage that one person can manage. The average space utilization in the open server world is 30%, vs. 80% for the mainframe world. CIM (Common Information Model) is a goal of a consortium of storage vendors, SNIA (Storage Network Industry Assn.), to develop a common interface for storage management software by year 2005. One of the big thrusts in storage management is storage virtualization, which allows one person to manage much more storage by breaking the relationship between the server operating system and physical storage. Storage virtualization may be implemented in four ways: asymmetric or metadata server, symmetric with a SAN storage manager, a switch-based system, and within the storage system. The storage management software field is not mature yet – with average software costs of $100K in an average environment of 40 servers.

b. BMC Users Group meeting on 9 Dec. 2002.

1. The day's agenda consisted of a combination of product direction statements and technical presentations. The morning portion was a general session with the afternoon activities conducted in parallel breakout tracks.

2. The keynote address was titled "Managing Enterprise Computing" by BMC President and CEO Bob Beauchamp. He noted that, even in these tough economic times, BMC is in a strong position, with zero debt and 1.23 Billion in cash and marketable securities. This year, BMC acquired the Remedy software product from Peregrine for $355 million in cash. The transaction was completed on 20 November. All seven senior Remedy executives have come on board with BMC. Dell has standardized on BMC software. Areas for potential growth include applications management (e.g. SAP) and LINUX. BMC intends to be a comprehensive enterprise management solution provider.

3. After the keynote, Richard Weiss of UPS presented a talk titled "Capacity and Performance Solutions at UPS". This presentation discussed the implementation details for BMC's Perform/Predict capacity management tools at UPS.

4. The UPS presentation concluded the morning session. After lunch, presentations were presented in four subject area tracks. One of these sessions was titled "Enterprise Performance Assurance Roadmap." This was presented by David Wagner, Director of Product Marketing for BMC. The mainframe and distributed systems development groups have been merged by BMC. A new product called Patrol Perceive is being added. This allows "ad hoc" consumers of performance data to gain access to the performance data via a web interface. Using pre-defined views, it allows users outside of the capacity management group (e.g. DBAs, Systems Administrators, Management) to "self-serve" capacity and performance data, thus leaving the capacity management group free to pursue higher value activities. Perceive is initially available only for UNIX and Windows platforms. This product could provide a low cost entry point for many activities just starting out in capacity management. BMC and IBM also announced that they were forming a partnership to develop and make available a free web-based toolkit to support CPU sizing for server consolidations involving LINUX on mainframes. No modeling is involved. Sizing estimates are provided by IBM.

5. Michael Ramos of BMC presented a talk titled “Patrol Network Management”. BMC is working on a product called Patrol Visualis (based on a product under development by a French company that BMC acquired). Visualis is a network and application traffic flow analysis application which is capable of modeling switched network topologies (layer 2) and is aware of Patrol agents present within the infrastructure. It also provides real time navigation and playback visualization for troubleshooting purposes. Traffic flow analysis can be done in navigation, VCR, or real time modes. Some of the features are similar to those found in Concord Network Health (a product used in our NeMO) but the Visualis product is less flexible, providing a basic, static functionality.

6. Lisa Morgan of BMC presented a talk titled “SLM Express & Patrol End to End Response timer. This included a description of the End to End Response Timer, a Mercury Topaz like product that can capture and replay transactions from robot replay stations. Response time data is then captured for the replay transactions.

7. On this day, BMC and SAS announced (at their respective user group meetings) the formation of a partnership to develop an integration module which will serve to combine BMC’s data collection and modeling functionality with SAS’s analytical reporting and chargeback capabilities. This is in development now.

c. Demandtech User Group on 9 Dec. 2002.

Demandtech held an all day user group meeting with a focus on software for performance analysis in the Windows world. Windows .Net is not in production yet for a server version. .Net offers a nice development environment under a GUI interface and it enforces 80% of software engineering best practices. It includes four command-line compilers: one for VB.NET, one for a new language called C# (C-Sharp), a new version of the C++ compiler and linker that produces what is called managed C++ (or MC++), and one for the Intermediate Language (IL). Another part of .NET is a development environment called Visual Studio.NET. .NET is primarily an architecture for writing applications that are object-oriented in nature, and both hardware and operating system agnostic. Version 5.2 of Windows is .Net server, which is in “release candidate status” instead of being in production status (with the holdup having been some problems porting it to Sun Solaris). The final rollout of version 5.2 has also been slowed by some concerns about licensing and security.

Below is a comparison of the .NET architecture compared to J2EE:

Technology	.NET	J2EE
Support technologies		
Distribution protocol	DCOM, SOAP	RMI/IIOP
Firewall	ISA*	not specified
HTML page caching	ISA*, ASP.NET	not specified
Presentation tier technologies		
Infrastructure	IIS	not specified
Programming model	ASP.NET	Servlets, JSP
High availability	NLBS*, ACS*, other	not specified
Load balancing	NLBS*, ACS*, other	not specified
Management	ACS*	not specified
Middle tier technologies		
Infrastructure	COM+	EJB
Programming tool	Visual Studio.NET	not specified
High availability	ACS*	not specified
Load balancing	ACS*	not specified
Security API	COM+ Security Call Context	JAAS
Message Queue API	MSMQ	JMS 1.0
Asynchronous components	Queued (COM+)	Message driven beans (EJB 2.0)
Naming and Directory Service	ADSI	JNDI
Data tier technologies		
Distributed transaction	MS-DTC	JTS
Relational DB API	ADO.NET	JDBC 2.0
Hierarchical DB API	ADO.NET	-
Database storage	SQLServer or ADO.NET compliant database	-
Mainframe DB connectivity	HIS*	Java Connectors
Framework technologies		
eCommerce framework	Commerce Server*	-
Business to business orchestration	BizTalk Server*	-

* Optional add-ons services for the .NET platform.

The .NET platform has a three year lead over J2EE in maturity.

d. SAS Users Group meeting on 9 Dec. 2002.

1. SAS held an afternoon user group meeting with the agenda consisting of a combination of product direction statements and technical presentations.

2. Jerry Carvalho of Morgan-Stanley presented a talk titled "An IT Resource Management View". At Morgan Stanley, they have implemented the SAS IT Resource Management (ITRM). SAS ITRM is an enterprise wide IT management solution that addresses a company's entire framework of services, systems, networks, phone systems and web servers. Warehousing and analyzing data provides reporting at all levels of your IT organization. Morgan-Stanley uses ITRM to take raw data staged from 4 regional servers located in New York, London, Tokyo and Hong Kong. The data is warehoused, analyzed and provides the capability for management at all levels to view the information through ITRM's reporting features.

3. Tom Lawler of IBM Global Services presented a talk titled "Optimizing the Value of Your Investment". Mr. Lawler discussed the value of IBM Global Services and SAS's IT Management Solutions approach together to optimizing your IT investment. IBM's performance management, testing and scalability services builds a solid infrastructure to support business objectives. They provide services that implement and test system, network, storage and application solutions for your environment through an automated integrated methodology based on the SAS IT Resource Management (ITRM) to improve accuracy and efficiency while performing tests and analyzing results.

e. CMG2002 International Conference 10-13 Dec. 2002.

1. Gary Bloom, the Chairman, President, and CEO of Veritas Software, gave the conference keynote address, titled "Architecting the Future of Storage". Gary sees the following trends in IT: server consolidation, vendor consolidation, and organizations pursuing the goals of reducing the complexity and cost of their IT infrastructures in order to improve Total Cost of Ownership (TCO). Today, there is an explosion of data production. Ironically, however, the actual storage hardware tends to be underutilized. Typically, for every dollar spent on storage hardware, three dollars are spent on storage management. Environments are getting very complex. Multiple vendors, explosive growth and the growing convergence of storage with the network all contribute to the cost of storage management. In the typical installation, labor accounts for about 58 percent of the cost of storage. Software can help make more efficient use of the hardware while reducing the cost of labor via automation. Many, many Fortune 500 companies are in the process of re-centralizing IT. Some are centralizing standards definition and dictating commonality, while others are doing actual physical consolidations. In the future, we will see the virtualization of storage with management intelligence moved outboard to a "switch" used to manage a virtual storage pool.

2. Robert S. Gold of Balanced Scorecard Collaborative, Inc. presented a paper titled "Enabling the Strategy-Focused Organization With Balanced Scorecard". The

CIO/CEO/CFO all ask "what is the value of IT." They look at it as an expense. Business unit managers ask "why can't you run IT like a business and focus on me, the customer". The CIO asks "how can IT meet the growing needs of the business units given the constraints on resources". IT needs to be seen as an information capital asset rather than an expense. Instead of concentrating on expense reduction, IT should concentrate on providing the greatest value. IT is an investment and should be managed via investment portfolios. The traditional IT focus on unit cost reduction and quality improvement is no longer enough. If IT cannot get out of a defensive mode, it will be outsourced. IT needs to speak the language of strategy. Balanced scorecard for IT helps IT to think strategically and creates a shared vocabulary with business. Less than 100 organizations in the nation are using balanced scorecard effectively, but they have a high success rate.

3. Pierre M. Fiorini of the University of Southern Maine presented a paper titled "On the Performability of Computing Systems: An Analytic Approach". Performability has its origins with John Meyer at the University of Michigan in the late 1970s. Performability modeling arose from a desire to evaluate computing systems that exhibit degradable performance. That is, they continue to operate despite the presence of failures in one or more of their components. Performability is a coined term combining performance and dependability. Performability models address both system performance and availability. The rest of this paper involves a quite technical discussion of mathematical methods. These methods can be useful in the study of distributed systems which include a great many components and can be designed to continue to function in a degraded mode even in the face of multiple failures.

4. Linwood Merritt of Capital One Services, Inc. presented a titled "A Capacity Planning Partnership with the Business". At Capital One, they have implemented Capacity Councils. These have a cross-departmental structure and a charter to work with the major business areas, technology organizations and support organizations to ensure that appropriate processing capability exists to support IT business customers. Each major business area has a capacity council. Each council has a support manager sponsor, is chaired by the business area and includes members from the capacity planning group, the business area and system/data administrators. Councils meet monthly and are responsible for reporting, collection and understanding of business drivers, capacity issues and major project and business plans. Application representatives develop estimates of business driver demand. Capacity planners translate this into resource level demand and capacity requirements. The goal of creating the councils was to develop a repeatable process for capacity management. Through their membership, the councils bring business and application knowledge into to the capacity planning process and create a partnership between the business entities, the application owners and the capacity planners.

5. Chris Malloy moderated a panel discussion titled "The Role of IT Metrics in IT Strategy". Panel members were Tom Bell, Bob Gold and Sidney Finehirsh. Sidney remarked that the utility model of IT is insufficient. It just makes IT invisible. IT needs to be a venture partner with the business. Bob noted that unit costs should be understood. Also, demand for

IT needs to be understood. They should be measured separately. All of the panelists agreed that there is a difference between tactical and strategic measures.

6. Fred Moore, President of Horison Information Strategies, presented a paper titled "Storage Manifesto". Fred is considered to be one of the foremost "gurus" in the storage arena. UNIX and Windows systems account for over 80 percent of all the digital storage installed. While storage capacity is scaling up fast, the performance of that storage is not keeping pace. Also, our ability to manage all of that storage can't keep up with the growth in storage quantity. So far, there is no real answer for the interoperability issues between storage vendors and there is an ongoing shortage of skilled personnel to perform storage management. In general, a small amount of data generates nearly 60 percent of all I/O. Solid state disk is making a comeback as a database accelerator. Today's disk HDAs are up to 180 GB. However, they tend to be only 30 to 40 percent utilized in order to keep I/O performance acceptable. In the long term (10+ years out), the vendors will be moving to remove all mechanical motion from disk devices. In the tape world, the renewed focus on disaster recovery in the wake of 9/11 is driving a large part of the market growth. While enterprise level disk subsystems still cost $70 - $90 per GB, automated tape libraries only cost $1.25 to $4.50 per GB. So, for cost reasons, tape is still a viable choice. In 10 years or so, tape cartridge capacity should reach the 4 – 8 TB range. Any disk volume made today can be backed up on a single tape cartridge. The "killer app" for tape is backup/recovery and Hierarchical Storage Management (HSM). So far, SANs have not been able to cope with the heterogeneous environments that most organizations have. In the average UNIX/LINUX/NT shop a storage administrator can only manage .75 TB. Over time, as capacity continues to grow, this will present a very real problem. You won't be able to hire enough administrators. The technology must improve. Optical disk has faded away and is no longer part of the storage hierarchy. In the nineties, there was an orgy of fiber optic installations. The existing fiber cable in the ground is underutilized. Much of it is idle (dark). This provides an opportunity for storage networking. However, the "last mile" problem does not have a real solution yet.

7. Jon William Toigo of Toigo Partners International presented a two part presentation titled "The Holy Grail of Data Disaster Avoidance: The Quest for Storage Recovery Strategies in a Nasty and Brutish World". Business processes have become inseparable from the technology infrastructure. Natural disasters are not more common, but are more costly. The essence of disaster recovery planning is an effort to eliminate avoidable disasters and to minimize the consequences of disasters that cannot be eliminated. It is not a one time project, but a process that is continuous across time. Traditionally, these planning tasks have been assigned to low level IT personnel. The traditional planning cycle involves an analysis phase (where critical applications and data are identified, data loss tolerance levels determined and restoration priorities assigned); a design phase (where the recovery strategies are developed); and an implementation phase (where the disaster recovery strategies are implemented and tested). IT evolution has presented challenges to this traditional planning cycle. PC proliferation, the advent of client/server distributed systems, the coming of the internet and the introduction of various networked storage schemes all introduce new and greater risk factors. However, at bottom, the primacy of data is unchanged. At the heart of all disaster recovery is the need to restore and recover your data. Without the data, your

applications are useless. While networked storage is not a panacea, it does offer some new possibilities for disaster recovery. In the final analysis, disaster recovery is merging with high availability. Data is the most irreplaceable asset and needs to be protected first. A common-sense data protection strategy is both proactive (detection and prevention) and reactive (post-disaster restoration). Networked storage solutions offer intriguing possibilities (e.g. data replication, etc.), but industry infighting over standards is hindering its usefulness in the heterogeneous computing environments common today. To date, tape remains the cheapest alternative and continues to be the workhorse in disaster recovery.

8. Greg Caliri of BMC Software, Inc. presented a paper titled "Outsourcing – what will happen and what will happen to YOU?" Outsourcing has existed in some form since the start of commercial computing. However, a fairly recent development has been the complete, or nearly complete, outsourcing deal (where an independent firm assumes nearly all IT responsibilities and the firm transfers IT personnel to the outsourcer). These days, outsourcing is always under consideration, even if an outsourcing decision is not imminent. For the technicians, most will either leave the organization or go with the outsourcer. The few left behind will be involved in Service Level Agreement observation and in liaison positions. For those technical types left behind, their knowledge base needs to be maintained. Conference attendance (e.g. CMG, SHARE), education, vendor sessions, etc. should continue. This knowledge is required in order to do proper oversight and liaison. The outsourcing agreement should clearly define that technicians remaining with the outsourcing organization have access to monitoring systems and records so that they can properly oversee and audit the outsourcer.

9. Fred Moore of Horisons, Inc. hosted a panel discussion titled "The Future of Storage Networking". Panelists included Mark Friedman of Datacore, Randy Kerns of Evaluator Group, Paul Massiglia of Veritas, Rob Pegler of Seagate, Brian Wong of Sun, Bill Zahlavi of EMC and Jon William Toigo. Subjects discussed included approaches to managing large scale storage networks, trends in storage virtualization, the use of IP networks in storage implementations and new technology trends. Some lively discussions ensued, especially in the area of standards, which have yet to emerge in any really useful form.

10. Al Sherkow of I/S Management Strategies, Ltd presented a paper titled "Mainframe Software Capacity Planning – IBM's Workload License Charges and Its Effect on Capacity Planning". IBM's move to Workload License Charges, use of defined capacity, soft caps on capacity, etc changes how capacity planning needs to be done for these environments. The environment becomes much more complex. Due to software pricing features such as variable pricing and capacity caps, how a system is sized and configured can affect the software bill drastically, sometimes driving it to be larger than the hardware bill. Capacity planners really need to study how LPAR configuration decisions affect the cost of software. All planning should be at the LPAR level. In this new environment, proper capacity planning can end up saving money on software charges. Mistakes can lead to unexpected expenses.

11. Adam Grummitt of Metron presented a paper titled "Creative Statistics: Are Performance Numbers Absolute or is Spin Inevitable in Their Presentation". This paper is a

discussion of how the choice of statistical analysis technique or graphic presentation method can create false impressions of what the data means. Regularly produced reports should have a consistent format and follow a selected "house style".

12. Dr. Bernie Domanski of The City University of New York, presented a paper titled "Software That Can Think and Do: Capacity on Demand Across Platforms". Rapid advances in research and technology allow data analysis and modeling of very complex systems. Methods from artificial intelligence such as neural networks have been explored to determine how they apply to systems management. It is possible to automate software systems to analyze measurements being generated in real time, automatically model its different facets, and then initiate actions automatically to maintain effective service time and resource utilization. By combining emerging technologies from AI, performance measurement and workload modeling, data centers will be able to begin enjoying a new level of availability and guaranteed service levels. Any look at total cost of ownership shows that the biggest cost is people. We need hardware that is able to manage itself much more than it does today. IBM has pulled together a number of efforts into one grand, multi-billion dollar research and development project (called eLiza) will the goal of creating self-managing, self-healing systems. IBM is working on scalable cluster management (project name Oceano), a highly integrated, parallel system that monitors everything from computing and network usage to application and database performance. It will automatically allocate computing resources to workloads as needed. Oceano is intended to introduce high levels of automation to dynamically adjust web sites to actual traffic demands over a massively parallel array of shared and distributed LINUX servers. Paul Horn, the Director of IBM research says that they have a prototype of this operating. He has said, "Building an autonomic computing network is <u>not</u> optional; in fact, the future of the Internet and e-business depends on it".

13. John Pilch of Performance Capacity Solutions presented a paper titled "Capacity Planning in the Development Cycle". Capacity planning can provide valuable contributions to the development process. It can help to determine economic feasibility by estimating what the cost per unit of work will be in production. It can contribute to the development strategy by helping to analyze options for server placement, use placement and rollout. It can contribute to configuration planning by evaluating existing configurations and how they will be affected by the new application. It contributes to capacity management for the new application by determining metrics, establishing instrumentation and identifying key indicators. Capacity planning activities require a reliable and consistent source of data collection, reporting and accurate tools to perform projections. This may require the application under development to license and fund software agents or tools. Measurement requirements must be developed during the business analysis phase and implemented during the development phase. Initially, measurement is required to develop the resource demands to be applied to the sizing model. In development, measurement data confirms the assumptions used in earlier phases and verifies the service levels achievable. Once in production, measurement is used for continuous performance monitoring, analysis of performance issues, identification of tuning opportunities, and building a database for ongoing capacity planning.

14. Dr. Connie U. Smith of Performance Engineering Services and Dr. Lloyd G. Williams of Software Engineering Research presented a paper titled "New Software Performance AntiPatterns: More Ways to Shoot Yourself in the Foot". Patterns and antipatterns are a hot topic in the software development world. A design pattern is a description of a general solution to a recurring design problem that can be customized to solve a particular problem. Design patterns capture expert knowledge about best practices in software design in a form that allows the knowledge to be reused. Antipatterns are similar to patterns in that they document recurring solutions to common design problems. They are known as antipatterns because their use (or misuse) produces negative consequences. Antipatterns document common mistakes and their solutions. Dr Smith and Dr Williams have identified performance antipatterns (common mistakes that lead to poor software performance). The authors then went on to review five antipatterns discussed at last year's CMG conference and four new antipatterns identified since then. The rest of the paper goes on to describe these newly identified antipatterns, their manifestations and their solutions.

15. Dr. Michael Salburg of UNISYS spoke on "The Thoroughly Modern Mainframe." He believes that large scale WINTEL (Window OS with an Intel chip) servers with up to 64 CPU's running one image will take over much of the server market. The business argument for such platforms is that on TPC benchmark tests, large scale Windows servers have a much lower cost per transaction. A problem in scaling such machines up beyond 8 machines is a performance penalty because of memory latency over the system bus. One way of addressing this problem is NUMA (Non-Uniform Memory Access), a type of parallel processing architecture in which each processor has its own local memory but can also access memory owned by other processors. It's called non-uniform because the memory access times are faster when a processor accesses its own memory than when it borrows memory from another processor. I/O speeds on Wintel servers are getting much faster: 133 Mhz bus speed now, 266 Mhz bus speed in 2003, and 5 Ghz speed projected for 2005.

16. Sanjay Raja of Caw Networks Inc. spoke on "SSL Solutions and Performance Criteria." SSL (Secure Sockets Layer) is a layer 4 protocol sitting between TCP and HTTP. SSL is a more lightweight protocol than IPSEC. SSL version 3 will let a session key live for up to 24 hours, while SSL version 2 lets it live for only a few minutes – there is a tradeoff between performance and security in deciding how long to let the key "live." Existing web servers are unable to support web traffic secured with SSL due to the burden that SSL processing puts on the CPU, so two approaches are used to address this: 1) add an encryption co-processor board to the server, or 2) use a separate SSL proxy appliance. In buying solution #2 one has to be careful what they are buying as some vendors will sell a PC with an SSL network card added and call it an "SSL accelerator."

17. Jerry Rosenberg of SRM Associates gave one of the best papers, "Monitoring NT Performance." His recommendation in troubleshooting problems was to first check memory performance, second the disks, and third the network. The performance monitor that comes with Windows is called "Perfmon" in Windows NT or known as "sysmon" in Windows 2000 (W2K). The monitor was built as a lab tool, not a performance management

tool and it can't save the data it records in a database. One of the weaknesses of Windows compared to mainframes is that Windows does not have a workload manager to manage multiple workloads running in one system image. For disk tuning, avoid disk compression and encryption. Windows NT doesn't measure disk busy, instead it measures queue length; W2K can measure disk busy. In network tuning, the server should have a faster NIC card than the clients. Also, network adapters that are not needed should not be loaded. In taking performance measurements, Perfmon can be run on a machine remote to the one you are measuring to minimize the effect of the Heisenberg uncertainty principle.

18. Scott Matsumoto of Xtremesoft, Inc. spoke on "Measuring the Performance and Scalability of Business Logic Components." Today many n-tier client server applications are built using components which are like software Legos, using either the COM+ model from Microsoft or J2EE from the Java model. This talk looked mainly at the COM+ model and found some problems in trying to do performance tuning on them. One problem is that many COM+ modules are embedded in different processes which all have the same name, "dllhost.exe." It is harder to figure out which modules are causing performance problems if many process names are the same in the Windows performance display. The trick for getting additional performance stats on the what is going on within these processes is to use the COM+ API. The speaker recommended getting the Windows Software Developer's Kit and the Resource Kit to extract additional tools which are available for performance analysis. In designing a COM+ application, tradeoffs need to be made between application response time and ease of fault isolation. For, calls across processes make fault isolation easier but have higher overhead (i.e. causes worse response time) than calls within a process. The speaker suggests a complex visual diagram system called Menasce's/Almeida's Client/Server Interaction Diagrams to represent the COM+ component relationships in the application.

19. Mark Austrowick of Demand Technology presented "UNIX I/O Performance Measurement Methodologies Applied to Old and New Storage Technologies." This was another presentation of a paper that was the best one at CMGI, the Italian CMG. Studies of both open systems and MVS show that average utilization of medium to large-scale disk farms is 20% or less, with very large systems averaging only 8%. The vanilla *iostat* command used to get disk statistics from open systems only shows logical disk activity; this talk described additional tools get more detailed performance statistics. A significant improvement in open system disk performance has been achieved in the last few years by introducing I/O path balancing and path failover (functions available with mainframes since the late 80's.). With the introduction of RAID devices, the storage architecture (microcode features, processor power, cache dimensions and management algorithm, raid level used, etc.) are more important than the disk characteristics themselves (e.g., seek time) in evaluating performance. *RAID devices should be considered primarily a data integrity and data availability solution instead of a performance solution.* If RAID performance is a concern, one should configure the RAID group with disks belonging to multiple SCSI adapters (rather than a single one) to achieve the maximum I/O back end parallelism.

20. "7 Deadly Sins of IP Performance Management" was presented by Laura J. Knapp of IBM. Application usage by end users is very unpredictable in IP - what was valid last week may not be valid today. Sometimes it is appropriate to block users after a given number have logged onto an application in order to conserve existing resources. The tools used should allow you to determine sessions by applications or bytes by applications. They should allow one to determine the top 10 clients by number of bytes transferred. Buffers are critical components of any application. Running low or out of buffers on any system can cause immediate application failure, a domino effect on related resources and applications, or intermittent application oddities. A large amount of data must be monitored in this environment including the utilization of the CPU, links, memory, the output queues, buffer misses, broadcast volumes, and FECN (forward explicit congestion notification) /BECN (backward explicit congestion notification) in a frame relay network. If the source terminal in a communications circuit generates frequent FECN bits, it indicates that the available network bandwidth (at that time) is not as great as can be supported by the destination terminal. This can occur because of outdated or inadequate network infrastructure, heavy network traffic, high levels of line noise, or portions of the system going down.

21. Chris Malloy of IBM discussed "Grid Computing." Grid computing is the computing technology that allows disparate computers (a grid of computers) to use resources (CPU, memory, and disk) between systems, without having to be tightly coupled. The grid allows one to separate work into small enough units of work in order to be processed by multiple smaller computers simultaneously, using the CPU cycles of machines which are idle. Grid computing is better suited for batch processes that are multi-threaded. While most computer grids involve computer resources, data grids have also become popular (e.g., Napster). Microsoft® has added function to its most recent operating system versions (Windows ME® and Windows XP®) to enable data grids. A well known example of grid computing is SETI (Search for Extraterrestrial Intelligence), software from the Univ. of California at Berkeley which uses a screen saver that can go get a chunk of data from a telescope at Arecibo over the internet, analyze that data for signs of extraterrestrial intelligence, and then report the results back to them. Some drug companies are using grid computing for drug simulations that would otherwise require a supercomputer. The Globus Project is an open source software grid computing project with software available at http://www.globus.org .

22. SAS Institute had a vendor booth that included information about SAS Learning Edition, a software product costing $125 retail that provides a copy of SAS for the PC along with tutorial materials, including a book, *Getting Started with SAS Learning Edition*. The copy of SAS provided with it is a single-user license can only process up to 1000 data points, and it expires at the end of 2006. It seems to be a very cost effective solution for anyone to get additional training in SAS.

23. The Evaluator Group is a vendor of storage management information that also had a vendor booth at CMG. They have competitive intelligence on the storage vendors which enables one to make comparisons between them. This information would be very time

consuming and difficult for an individual to gather, especially as some storage vendors (e.g., EMC) include a standard clause in their contracts that you are not allowed to disclose performance data about their products. Some government agencies such as DISA have Evaluator Group information to save a lot of money.

24. Dr. Thomas Bell presented a paper titled 'Managing Computer Performance Consulting'. This paper is a discussion of some of the advantages of becoming a computer performance consultant and the challenges of managing the different aspects of the consulting world. Some of the areas Dr. Bell discussed involved the client management, time management, non-billable time management and the added administrative responsibilities of performing consultant work. Dr. Bell also discussed the ethical compromises involved when consultants want to increase sales or increase profits and provided suggestions for long term performance consulting.

25. Tony Shediak of IBM Global Services Australia presented "Performance Tuning Mainframe Applications". While DISOC is no longer involved in the day-to-day capacity management of the mainframe workload processing, this session offered numerous tidbits of performance improvement knowledge which might come in handy in the longer-term. Mr. Shediak discussed many performance tuning guidelines, areas that often contain problems, and some detailed methods of dealing with these problems. There are a number of very good suggestions for application areas to consider changing code to yield better performance 'from the inside'. He also mentioned the value of using a tool such as STROBE to monitor code for performance problems, which DISA offers DLA's legacy applications processing.

26. The panel session "Managing Performance People and Other Technical Personnel" was chaired by Dr. Tom Bell of Rivendel Consultants. Panelists included one of the DISOC attendees, with 3 other 4 panelists. Discussion centered on areas of concern to managers managing performance technicians, hiring issues in today's economy, ways to improve incentives to keep such personnel, and related topics.

27. Andrew Odlyzko of the University of Minnesota discussed "Measurements and Mismeasurements and the Dynamics of Data Traffic Growth". Dr. Odlyzko was an invited speaker who specializes in the collection of data used to measure Internet traffic. His research into the high rates of growth originally predicted for Internet traffic showed that contrary evidence was plentiful, had anyone paid attention. The interesting depiction of graphs and data was challenged by Dr. Neil Gunther, who was somewhat skeptical of the data, but Dr. Odlyzko skillfully showed that the data he presented was accurate. Dr. Gunther accepted the argument, which went a long way to adding credibility to the presentation. It was an absorbing discussion of the nuances of data collection, yielding the observation that the method of collection is key, and also that the data must be carefully interpreted.

28. "An Introduction to the InfiniBand Architecture" session was presented by Odysseas Pentakalos of SYSNET International, Inc. The InfiniBand architecture has been planned to provide the backbone for server I/O subsystems in which limitations of current bus architectures can be reduced or eliminated while also taking advantages of increases in CPU speeds, networking, and various newer storage configurations. Unfortunately, while originally many storage and related hardware vendors joined forces to promote this capability, most have since dropped out of the consortium. As it stands, the InfiniBand architecture is a highly-preferred concept technically, according to industry experts, but the implementation of this concept remains to be seen as to who (if anyone) stands behind it for the long haul.

29. "Linux and zSeries on their way to the eServer of the Future" was presented by Dr. Karl-Heinz Strassemeyer, IBM Corp. 'Dr. Karl', an IBM Distinguished Engineer, discussed his insights into the history and evolution of the large scale server architecture, as well as his outlook for the future of the IT world.

30. Bob Yellin of IBM/Tivoli Software spoke on "Predicting Service Level Assurance", discussing the need for standards-based service level management tools to measure whether or not Service Level Agreements are being met or not.

31. "Celebrity Boxing (and Sizing): Alan Greenspan vs Gene Amdahl!" was the title of the session by Neil Gunther of Performance Dynamics. This presentation centered on which of many sizing models one should use. He focused on two commonly-used models: Amdahl's Law and Greenspan's Law (of which, interestingly, neither famous person invented 'their' law). Conclusions after much esoteric mathematical discussion: Greenspan's Law is bogus for use with large-scale processor configurations; if throughput is not overdriven, then use Amdahl's Law; otherwise, the real winner is….(surprise) a sizing model previously presented by Dr. Gunther at CMG in the 1993-1996 timeframe.

32. "Workload Manager – Revisiting Goals over Time" was presented by Peter Enrico of Enterprise Performance Strategies, Inc. While DLA no longer has direct responsibility for mainframe processing performance, this session was a good time to update the knowledge of a key mainframe capability by a very good instructor. Mr. Enrico elaborated on the basic theory of WLM, and then progressed to concrete examples of various scenarios containing common poorly-defined goals. He presented several web sites where more information could be obtained, including www.sherkow.com.

5. CONCLUSIONS:

a. As usual, the CMG conference proved itself to be very valuable and included lots of useful information.

b. The conference agenda continues to provide very strong content in the areas of distributed computing, UNIX, Windows NT and, increasingly, electronic commerce. CMG continues to address areas of interest and concern for DLA.

6. RECOMMENDATIONS:

DISOC-CIIO should continue to attend the CMG conference and should continue to send multiple persons in order to cover the multiple simultaneous tracks.

James M. Haskins Kathy J. Steffens Mott Given Scott Engel